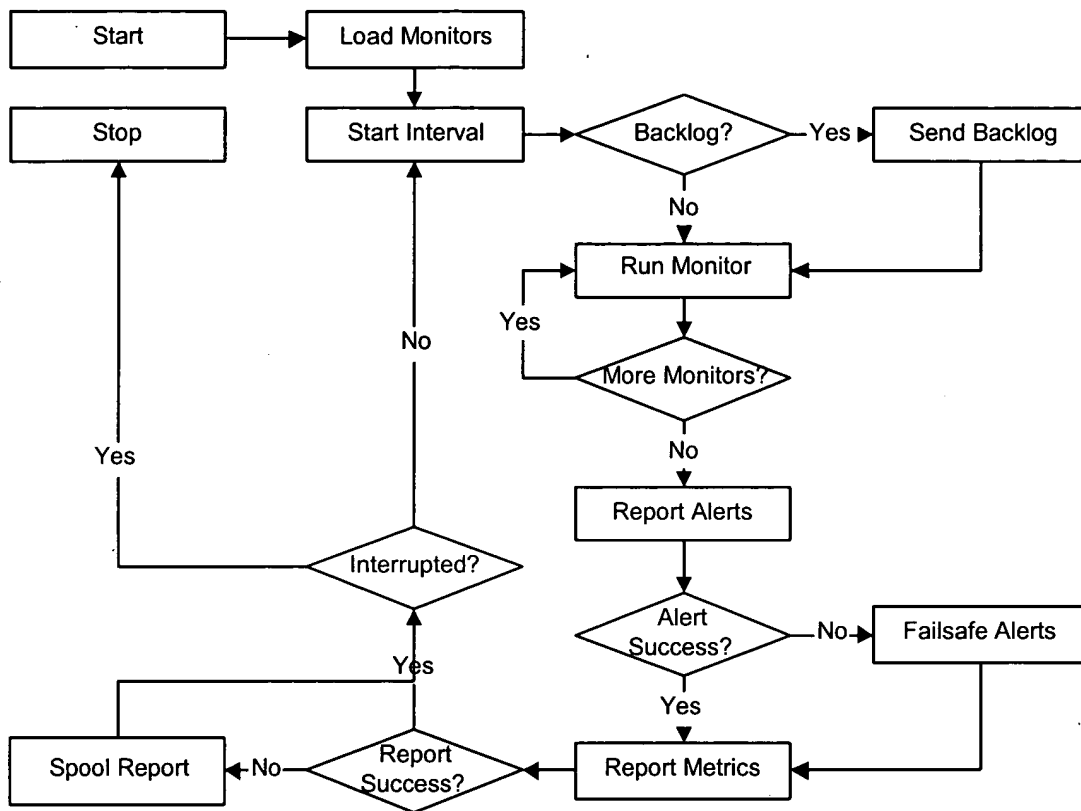


**FIG. 1**



**FIG. 2**

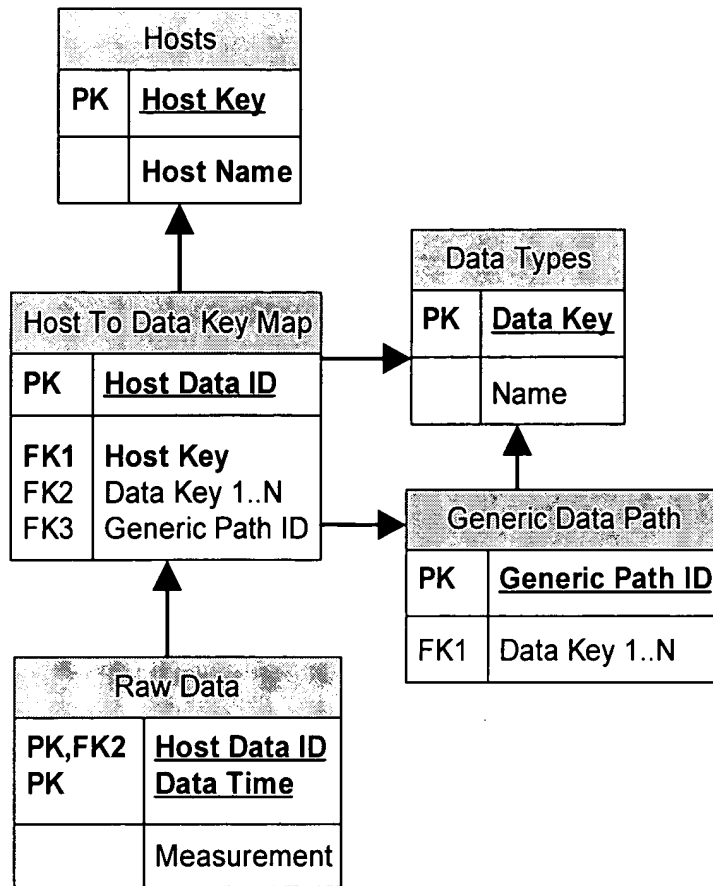
```

<?xml version="1.0"?>
<!DOCTYPE import [
<!ELEMENT import (client+)>
<!ELEMENT client (host+)>
<!ELEMENT host (ma*, log*, re*, li*)>
<!ELEMENT madata (m, s)>
<!ELEMENT m (#CDATA)>
<!ELEMENT s (#CDATA)>
<!ELEMENT log (m, s)>
<!ELEMENT re (m)>
<!ELEMENT li (m)>
<!ATTLIST client name CDATA #REQUIRED>
<!ATTLIST host name CDATA #REQUIRED>
<!ATTLIST host time CDATA #REQUIRED>
<!ATTLIST ma dk CDATA #REQUIRED>
<!ATTLIST ma time CDATA #IMPLIED>
<!ATTLIST log dk CDATA #REQUIRED>
<!ATTLIST log time CDATA #IMPLIED>
<!ATTLIST re type CDATA #REQUIRED>
<!ATTLIST re time CDATA #IMPLIED>
<!ATTLIST li time CDATA #IMPLIED>
<!ATTLIST li interval CDATA #REQUIRED>
]>

<import>
  <client name=company1>
    <host name="host.company1.com" time="1000000000">
      <ma dk="OS|HP-UX|CPU|pct_busy" time="1000000011">
        <m>98</m>
        <s>2</s>
      </ma>
      <log dk="OS|HP-UX|Log|Syslog">
        <m> ... </m>
        <s>2</s>
      </log>
      <re type="Disk Configuration">
        <m>Text of report</m>
      </re>
      <li dk="App|Lawson|614|Transactions|gl_trans|trans_day"
time="1000000000" interval=86400>
        <m>1000</m>
      </li>
    </host>
  </client>
</import>

```

**FIG. 3**



**FIG. 4**

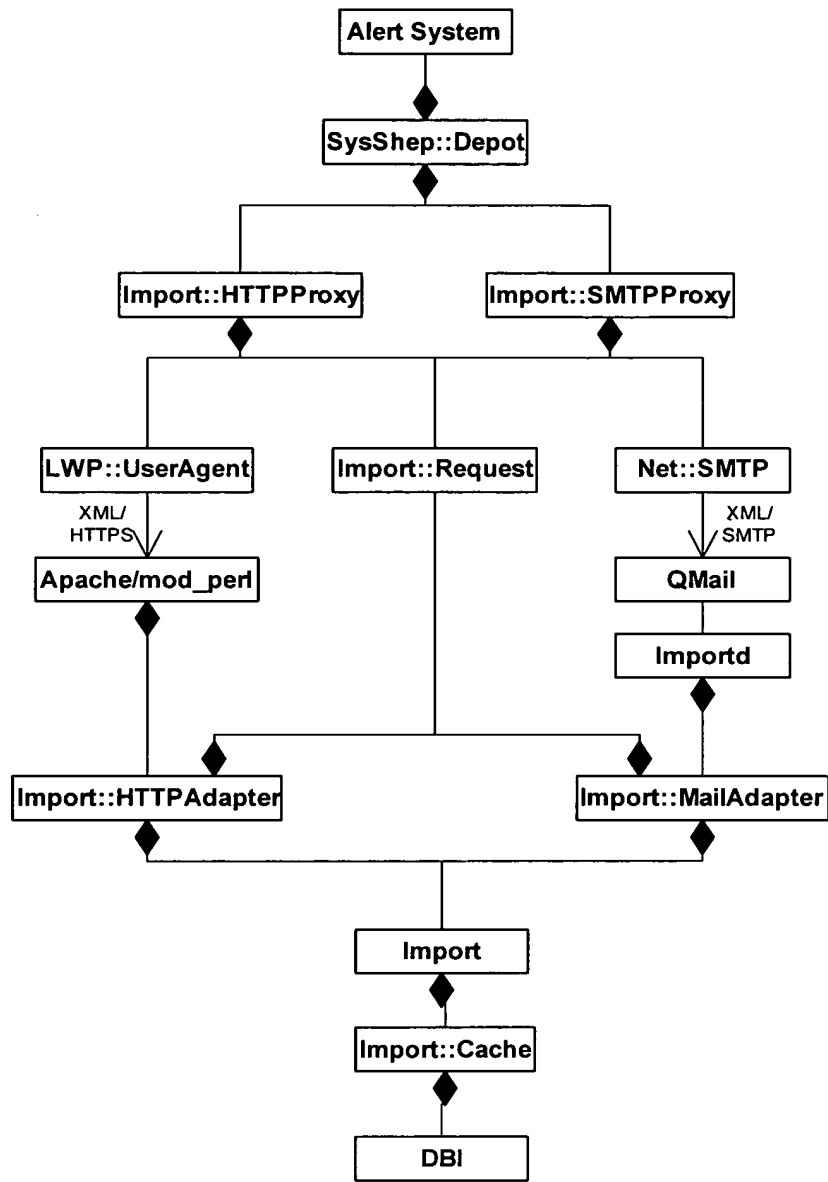
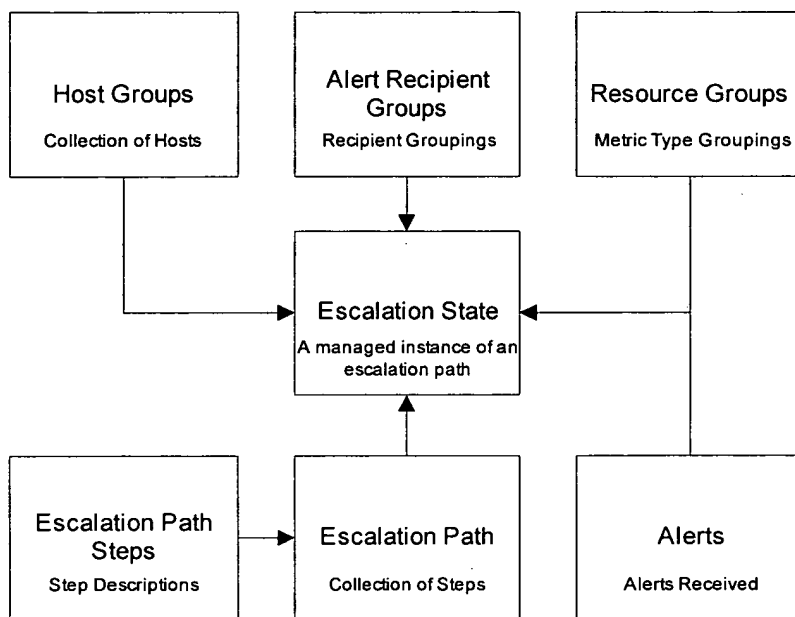
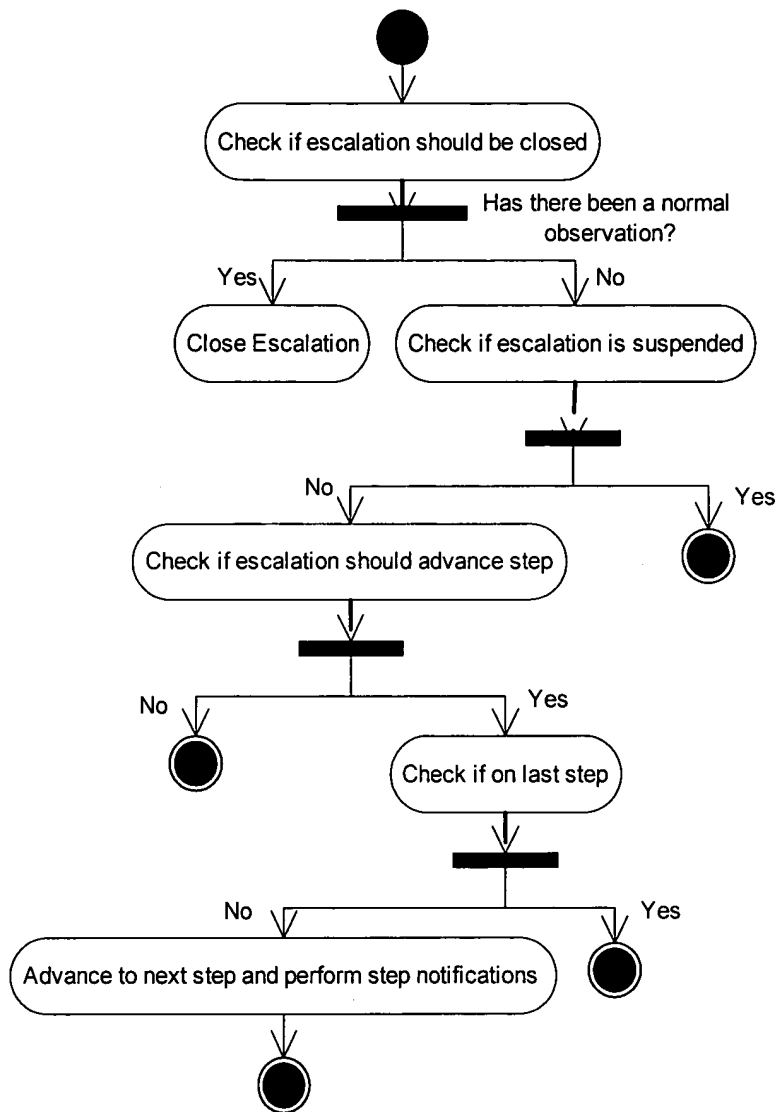


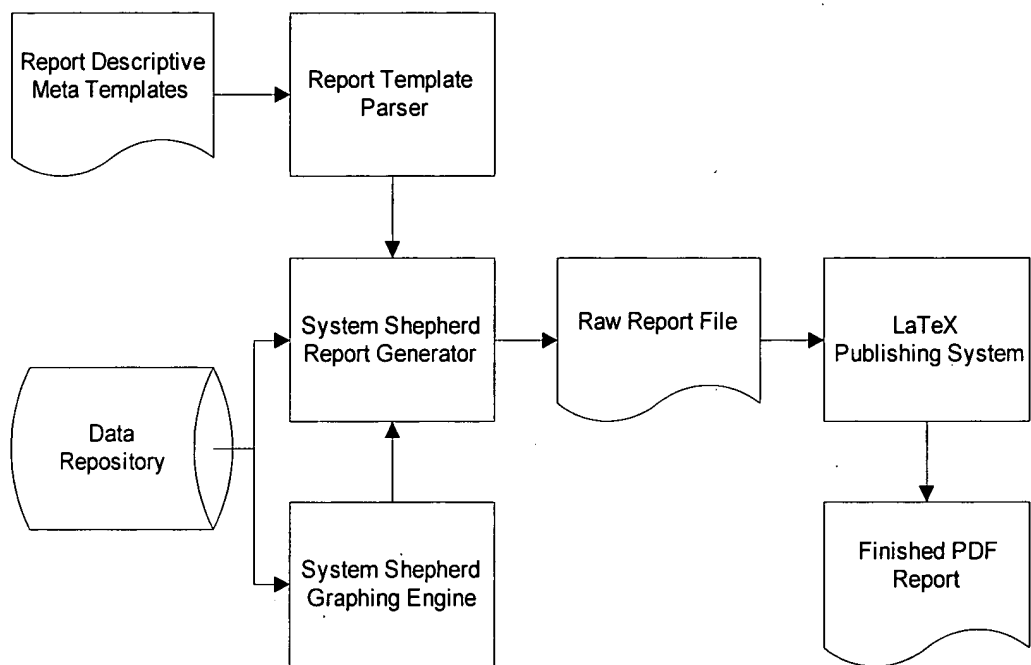
FIG. 5



**FIG. 6**



**FIG. 7**



**FIG. 8**



## 5.1 Total Latch Wait Time

The total latch wait time is the amount of time (in milliseconds) that all processes have waited for a latch, and is cumulative from the startup of SQL Server. The graph below can provide two insights into the performance of SQL Server:

- An indication that SQL Server was stopped and restarted (the counter was reset back to 0)
- Can illustrate periods of time that latch contention is highest (If the graph jumps at consistent periods of time, such as when nightly batch jobs execute)




Figure 10: Total Latch Wait Time in ms

### Analysis Results

The maximum Total Latch Wait Time is 714.30 ms and is above 30 ms. It is possible this problem is due to a disk bottleneck, where the length of the latch held is unusually long while the disk heads must seek to the proper location before reading the data. Another possibility is that the buffer cache is sized too small so that many queries must read their data from disk (instead of the buffer cache). It is a good idea to examine the buffer cache hit ratios, the disk queue lengths (from the OS report) and the number of latches/sec that are being requested.

**FIG. 9**

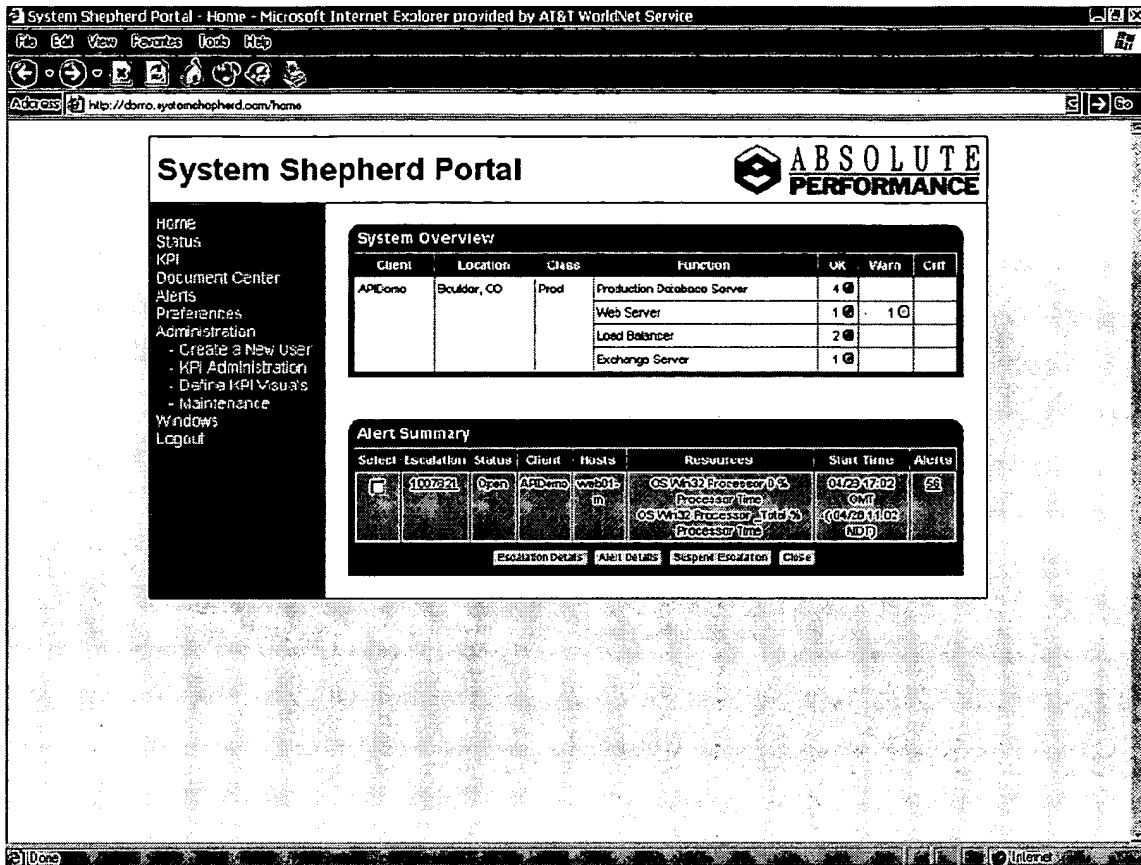


FIG. 10

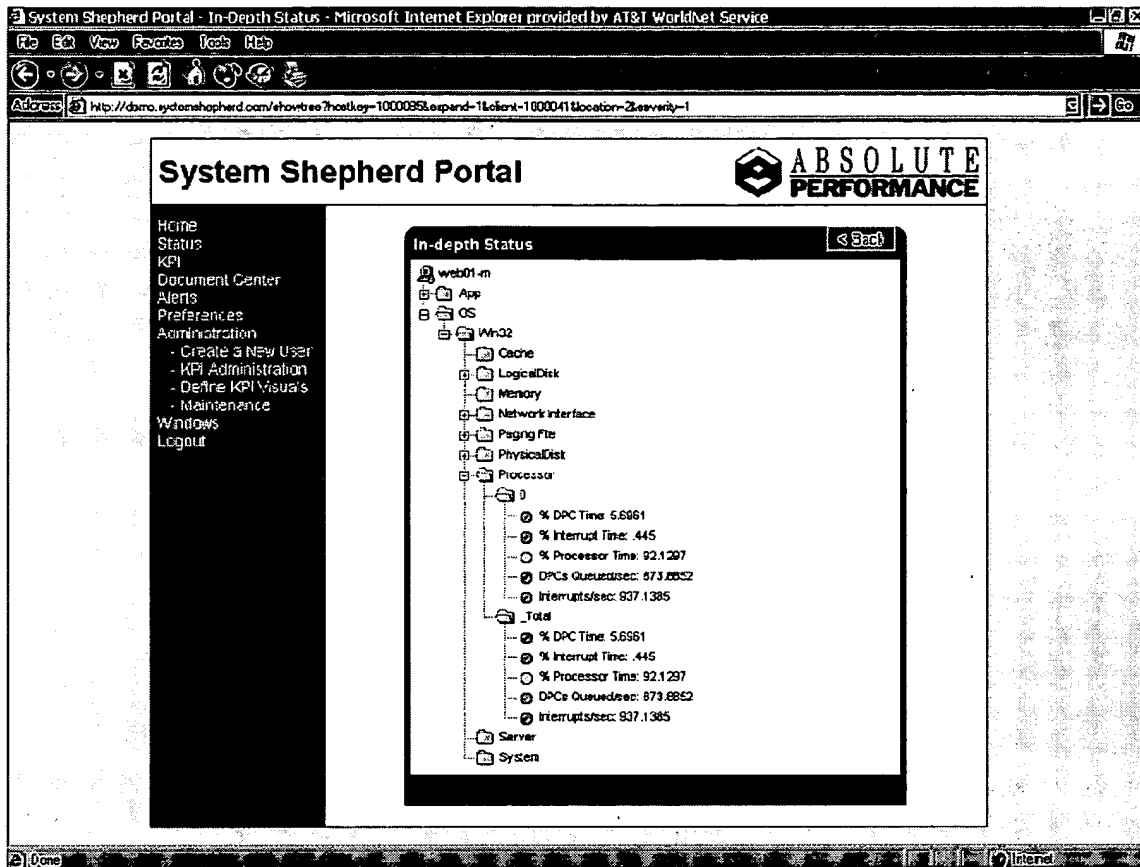


FIG. 11




FIG. 12

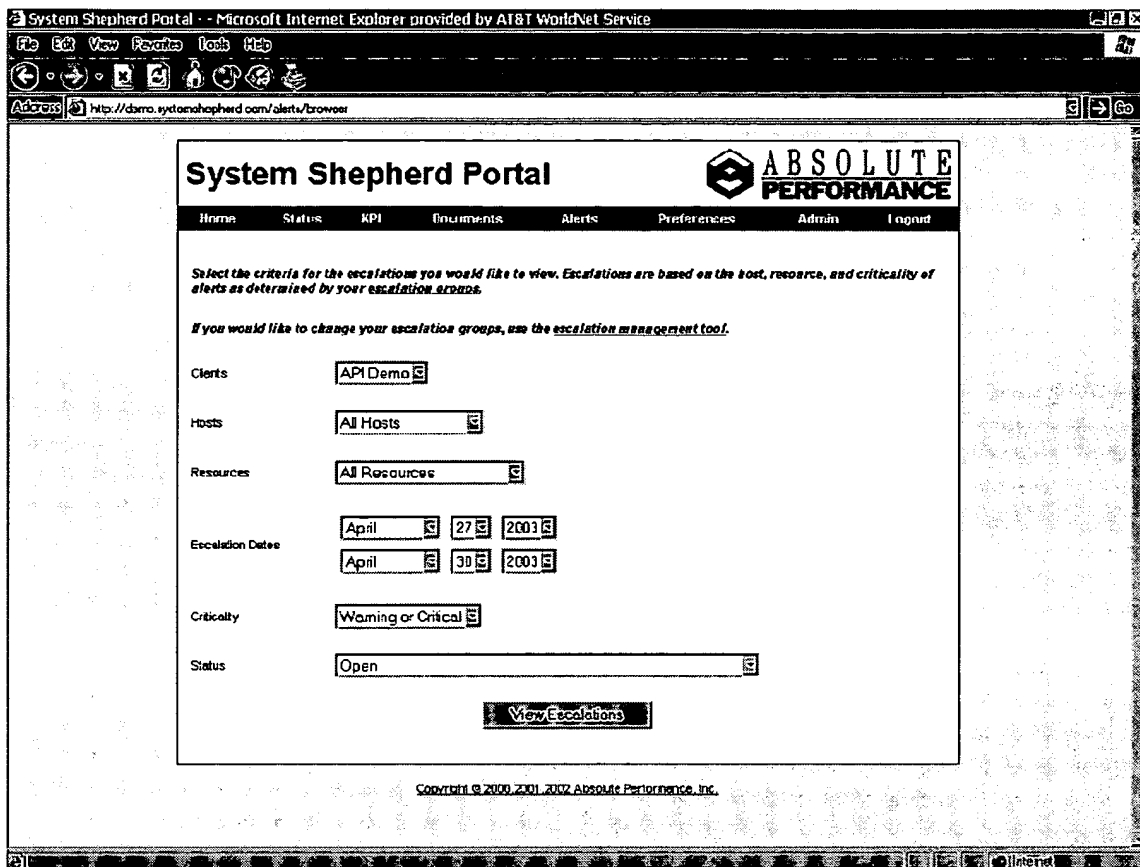


FIG. 13

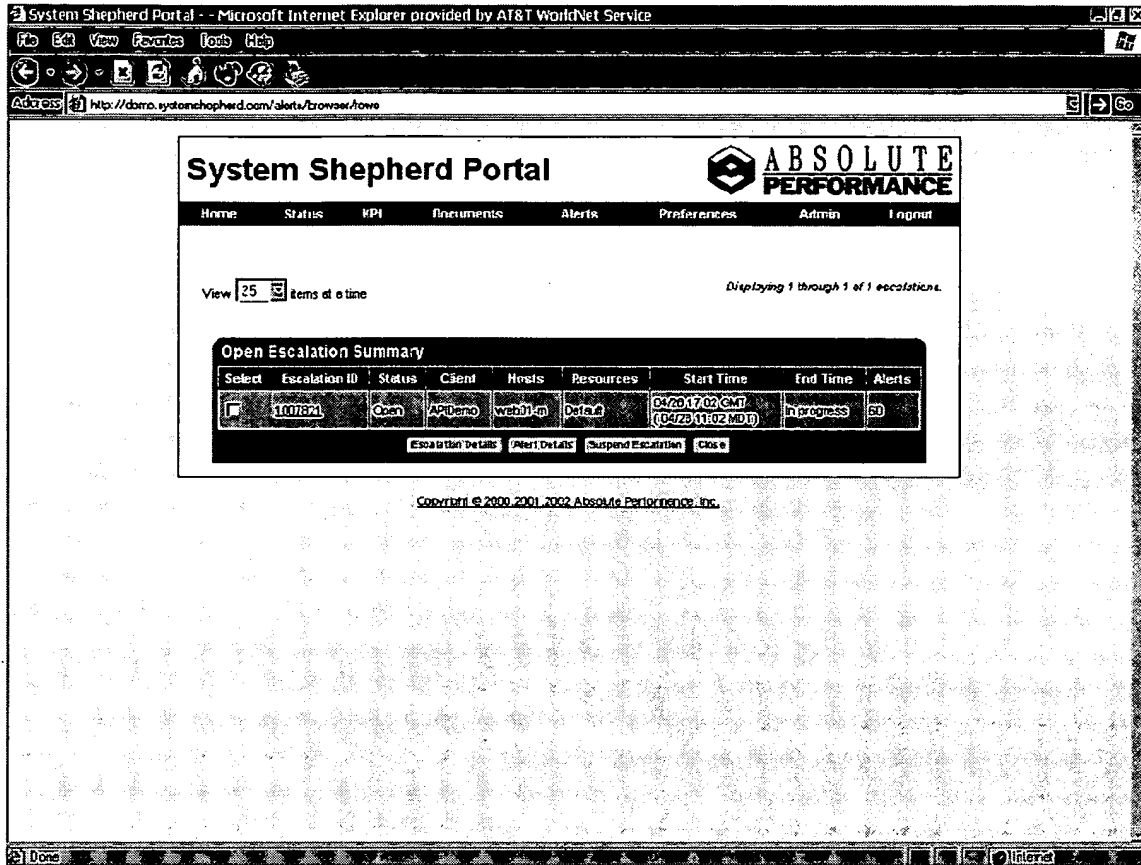


FIG. 14

System Shepherd Portal - Microsoft Internet Explorer provided by AT&T WorldNet Service

File Edit View Favorites Tools Help

Address http://demo.systemshepherd.com/alerts/browser/escalations

### System Shepherd Portal

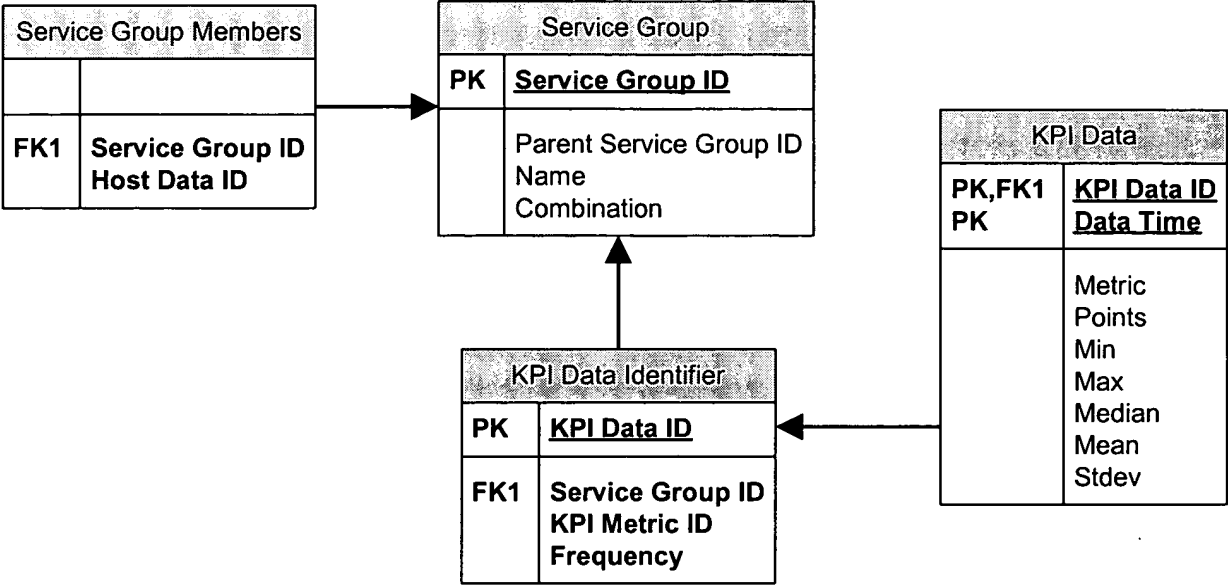
**ABSOLUTE PERFORMANCE**

Home Status KPI Documents Alerts Preferences Admin Logout

Escalation 1007821	
Escalation Group ID	1000083
Status	Open
Client	APDemo
Resource Group	Default
Alerting Hosts	web01-in
Number of Alerts	50
Start Time	2003-04-28 17:02:46 GMT ( 2003-04-28 11:02:46 MDT )
End Time	In progress
Current Escalation Step	1
Time of Next Escalation Step	2003-04-28 18:02:46 GMT ( 2003-04-28 12:02:46 MDT )
Last Alert Notification	2003-04-28 17:48:33 GMT ( 2003-04-28 11:48:33 MDT )
This escalation has not been suspended.	
This escalation has not been acknowledged.	

Copyright © 2000, 2001, 2002 Absolute Performance, Inc.

FIG. 15



**FIG. 16**



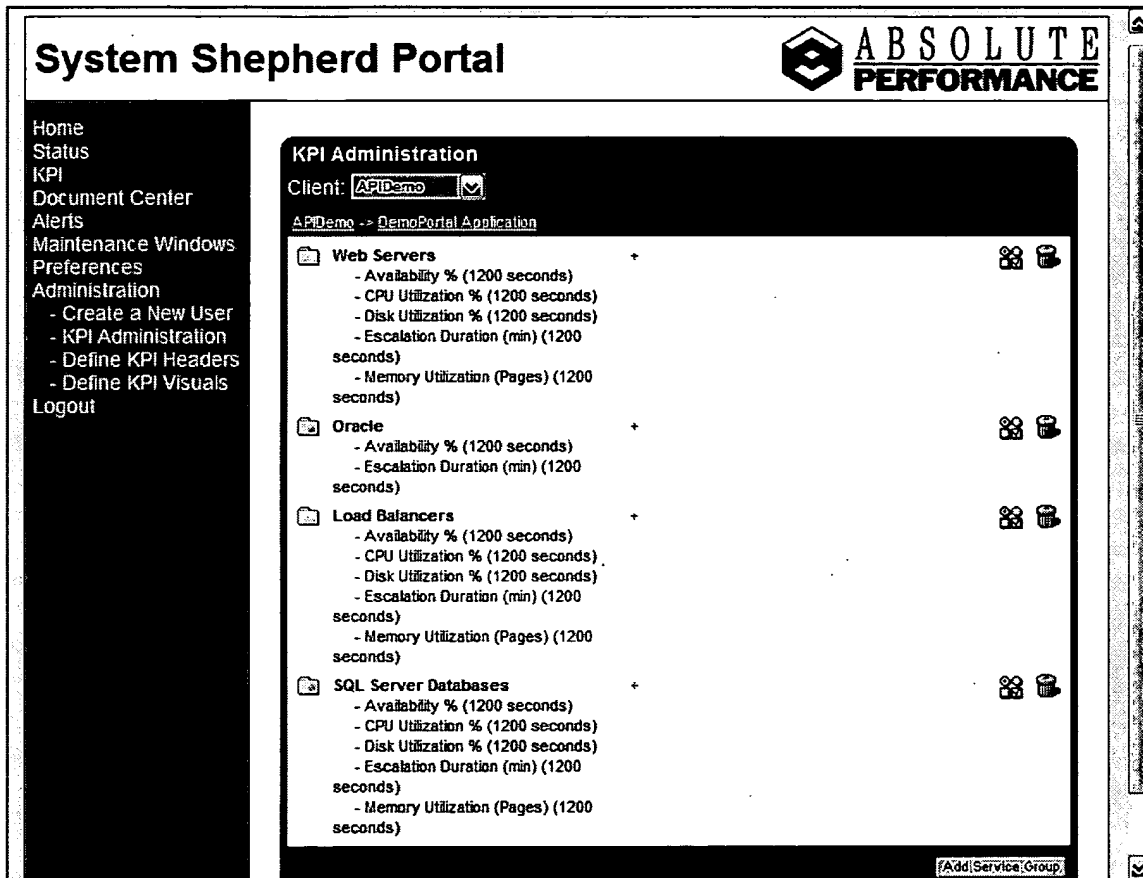


FIG. 17




FIG. 18